arXiv:1612.01205v1 [stat.ML] 4 Dec 2016

Optimal and Adaptive Off-policy Evaluation in Contextual Bandits

Yu-Xiang Wang[1], Alekh Agarwal[2], and Miroslav Dudík[2]

[1]Carnegie Mellon University, Pittsburgh, PA
[2]Microsoft Research, New York, NY

Abstract

We consider the problem of off-policy evaluation—estimating the value of a target policy using data collected by another policy—under the contextual bandit model. We establish a minimax lower bound on the mean squared error (MSE), and show that it is matched up to constant factors by the inverse propensity scoring (IPS) estimator. Since in the multi-armed bandit problem the IPS is suboptimal [Li et al., 2015], our result highlights the difficulty of the contextual setting with non-degenerate context distributions. We further consider improvements on this minimax MSE bound, given access to a reward model. We show that the existing doubly robust approach, which utilizes such a reward model, may continue to suffer from high variance even when the reward model is perfect. We propose a new estimator called SWITCH which more effectively uses the reward model and achieves a superior bias-variance tradeoff compared with prior work. We prove an upper bound on its MSE and demonstrate its benefits empirically on a diverse collection of datasets, often seeing orders of magnitude improvements over a number of baselines.

1 Introduction

Contextual bandits refer to a learning setting where the learner repeatedly observes a context, takes an action and observes a reward signal for the quality of the chosen action in the observed context. Crucially, there is no information on the quality of all the remaining actions that were not chosen for the context. As an example, consider online movie recommendation where the context describes information about a user, actions are possible movies to recommend and a reward can be whether the user enjoys the recommended movie. The framework applies equally well to several other applications such as online advertising, web search, personalized medical treatment, etc. The goal of the learner is to come up with a policy, that is a scheme for mapping contexts into actions. A common question which arises in such settings is, given a candidate *target policy*, what is the expected reward it obtains? A simple way of answering the question is by letting the policy choose actions (such as make movie recommendations to users), and compute the reward it obtains. Such *online evaluation*, is typically costly and time consuming since it involves exposing users to an untested experimental policy, and does not easily scale to evaluating the performance of many different policies.

Off-policy evaluation refers to an alternative paradigm for answering the same question. Suppose we have existing logs from the existing system, which might be choosing actions from a very different *logging policy* than the one we seek to evaluate. Can we estimate the expected reward of the *target policy*? This question has been extensively researched in the contextual bandit model (see, e.g., [Li et al., 2011, Bottou et al., 2013, Oh and Scheuren, 1983, Little and Rubin, 2002] and references therein). In particular, there are several estimators which are unbiased under mild assumptions, such as inverse propensity scoring (IPS) [Horvitz and Thompson, 1952], and sharp estimates on their mean squared error (MSE) for policy evaluation are well-known [Dudík et al., 2014].

While the IPS-style methods make no attempt at all to model the underlying dependence of rewards on contexts and actions, such information is often available. The simplest approach to off-policy evaluation, given such a model, is to simply use the model to predict the reward for the target policy's action on each context. We call this estimator the model-based approach or the direct method (DM). The key drawback of DM is that it can be arbitrarily biased when the model is misspecified. Some approaches, such as the doubly-robust method (DR) [Dudík et al., 2014] (also see the references therein for its origin in statistics and application in causal inference, e.g., [Robins and Rotnitzky, 1995, Bang and Robins, 2005]), combine the model with an IPS-style unbiased estimation and remain consistent, with known estimates for their MSE.

All these works focus on developing specific methods alongside upper bounds on their MSE. Little work, on the other hand, exists on the question of the fundamental statistical hardness of off-policy evaluation and the optimality (or the lack of) of the existing methods. A notable exception is the recent work of Li et al. [2015], who study off-policy evaluation in multi-armed bandits—a special case of our setting, without any contexts—and provide a minimax lower bound on the MSE. Their result shows the suboptimality of IPS (and DR) due to an excessive variance of the importance weights. This result is rather intriguing as it hints at one of two possibilities: (i) IPS and variants are also suboptimal for contextual bandit setting and we should develop better estimators, or (ii) the contextual bandit setting has qualitatively different upper and lower bounds that match. In this quest, our paper makes the following key contributions:

1. We provide the first rate-optimal lower bound on the MSE for off-policy evaluation in contextual bandits. In contrast with context-free multi-armed bandits [Li et al., 2015], our lower bound matches the MSE upper bound for IPS up to constants, so long as the contexts have a non-degenerate distribution. This highlights the challenges of the contextual setting; even if the reward as a function of contexts and actions has no variance, the lower bound stays non-trivial in contrast with context-free multi-armed bandits.

2. We propose a new class of estimators called the SWITCH estimators, that adaptively interpolate between an available reward model and IPS. We show that SWITCH has MSE no worse than IPS in the worst case, but is robust to large importance weights. We also show that SWITCH can have a drastically smaller variance than alternatives for combining IPS with a reward model, such as DR.

3. We conduct experiments showing that the new estimator performs significantly better than existing approaches on simulated contextual bandit problems using real-life multiclass classification data sets.

The rest of this paper is organized as follows. In Section 2 and 3, we recap the formal setup of off-policy evaluation and present a lower bound on the minimax risk for model-free off-policy evaluation in contextual bandits. In Section 4, we discuss the need to have adaptive estimators that can take advantage of a given reward model and present the SWITCH estimator along with its theoretical properties. Lastly, the experimental comparison of the proposed methods and existing approaches are given in Section 5.

2 Setup and related work

We use the standard setup for off-policy evaluation in contextual bandits as in prior works Dudík et al. [2014]. We start with the formal setup and then discuss the related work.

2.1 Formal setup

In contextual bandit problems, the learning agent observes a context x, takes an action a and observes a scalar reward measuring the quality of the chosen action for the context. Here the context x is a feature vector from a finite set $\mathcal{X}$, possibly $\{0,1\}^d$ or a machine-precision encoding of $\mathbb{R}^d$ for some large d. The stationary

distribution of contexts is denoted by λ. Actions, denoted as a, are drawn from a finite set $\mathcal{A}$. Rewards r have a distribution conditioned on x and a denoted by $D(r|x,a)$. The decision rule of the agent is called a policy, which is a function from contexts to distributions over actions to allow for randomized action choice. We will use $\mu(a|x)$ and $\pi(a|x)$ to denote the *logging* and *target* policies respectively. Given a policy π which is a distribution over actions given contexts, we extend it to a joint distribution over triples (x,a,r), where x is drawn according to λ, action a according to $\pi(a|x)$, and r according to $D(r|x,a)$. With this notation, the goal of off-policy evaluation is to evaluate the expected reward of π, given samples where actions are chosen according to μ. Formally, given n samples of the form (x_i,a_i,r_i,p_i), where $p_i=\mu(a_i \mid x_i)$, we wish to compute the *value of* π:

$$v^\pi = \mathbb{E}_\pi[r], \qquad \text{where we have } x\sim\lambda,\ a\sim\pi(\cdot\mid x), \text{ and } r\sim D(r\mid a,x). \tag{1}$$

In order to correct for the mismatch in the action distributions under μ and π, it is typical to use *importance weights*, and it will be convenient to introduce the shorthand $\rho(x,a)=\pi(a|x)/\mu(a|x)$. In order to facilitate consistent estimation, it is typical to assume that $\rho(x,a)\neq\infty$, meaning that whenever $\pi(a|x)>0$, then $\mu(a|x)>0$ as well. We will make this same assumption throughout our upper and lower bounds. We conclude the setup by giving one concrete example for estimating the quantity v^π to aid the reader's intuition. The simplest estimator in this setting, called Inverse Propensity Scoring (IPS) [Horvitz and Thompson, 1952] is defined as:

$$\hat{v}^\pi_{\mathrm{IPS}} = \sum_{i=1}^{n} \rho(x_i,a_i) r_i. \tag{2}$$

In the sequel, we will discuss properties of this estimator, as well as improvements upon it for off-policy evaluation.

2.2 Related work

We now review the related work. There are several works that focus on creating estimators for the off-policy evaluation problem, and provide upper bounds on the Mean Squared Error (MSE) of these estimators, as discussed in the introduction. The study of lower bounds on the MSE of *any possible estimator* for this problem is relatively new, on the other hand. Li et al. [2015] first studied this question under a minimax setup, but focused on the multi-arm bandits problem. As we will show, the naive extension of their results to contextual bandits yields a suboptimal lower bound. Jiang and Li [2016] built on top of Li et al. [2015] proves a Cramer-Rao-style lower bound for policy-evaluation in reinforcement learning, but the assumption of a small number of *observed states* precludes contextual bandits from being a special case. Cortes et al. [2010] studied the related covariate shift problem in the statistical learning setting and proved upper bounds and lower bounds of the minimax excess risk that depends on the Renyi divergence of the two covariate distributions. This is related to our problem in the following sense: we can formulate off-policy evaluation as a statistical learning problem of predicting reward r using feature vector (x,a), and off-policy evaluation corresponds to exactly the covariate shift problem (when we know how the covariate has shifted). However, as the minimax problems are defined differently[1], their results (especially the lower bound) are not applicable to our problem.

On the complementary question of developing techniques for off-policy evaluation and understanding their theoretical properties, there is a considerably larger body of work in both the statistics and machine learning literature. In statistics, this is studied as the problem of mean estimation for treatment regimes in the causal inference literatre (see e.g. [Horvitz and Thompson, 1952, Holland, 1986, Bang and Robins, 2005, Rotnitzky et al., 2012] as well as the references in [Dudík et al., 2014]). The "doubly robust" estimation techniques

[1] Off policy evaluation focuses on estimating the policy-value, while statistical learning focuses on prediction accuracy measured by the expected loss (in terms of the excess risk).

[Cassel et al., 1976, Robins and Rotnitzky, 1995], were recently used for off-policy value estimation in the contextual bandits problem [Dudík et al., 2011, 2014] and reinforcement learning [Jiang and Li, 2016]. These also provide error estimates for the proposed estimators, which we will compare with the lower bounds that we will obtain. The doubly robust techniques also have a flavor of incorporating existing reward models, an idea which we will also leverage in the development of the SWITCH estimator in Section 4. We note that similar ideas were also recently investigated in the context of reinforcement learning by Thomas and Brunskill [2016].

3 Limits of off-policy evaluation

In this section we will present our main result on lower bounds for off-policy evaluation. We first present the minimax framework to study such problems, before presenting the result and its implications.

3.1 Minimax framework

Off-policy evaluation is a statistical estimation problem, where the goal is to estimate v^π given n iid samples generated according to a policy μ. We study this problem in a standard minimax framework and seek to answer the following question. What is the smallest mean square error (MSE) that *any* estimator can achieve in the worst case over a large class of contextual bandit problems. As is usual in the minimax setting, we want the class of problems to be rich enough so that the estimation problem is not trivialized, and to be small enough so that the lower bounds are not driven by complete pathologies. In our problem, we make the choice to fix λ, μ and π, and only take worst case over a class of reward distributions. This allows the upper and lower bounds, as well as the estimators to adapt with and depend on λ, μ and π in interesting ways. The family of reward distributions $D(r \mid x, a)$ that we study is a natural generalization of the class studied by Li et al. [2015] for multi-armed bandits.

To formulate our class of reward distributions, assume we are given maps $R_{\max} : \mathcal{X} \times \mathcal{A} \to \mathbb{R}_+$ and $\sigma : \mathcal{X} \times \mathcal{A} \to \mathbb{R}_+$. The class of conditional distributions $\mathcal{R}(\sigma, R_{\max})$ is defined as

$$\mathcal{R}(\sigma, R_{\max}) \coloneqq \Big\{ D(r|x,a) : 0 \le \mathbb{E}_D[r|x,a] \le R_{\max}(x,a) \text{ and } \mathrm{Var}_D[r|x,a] \le \sigma^2(x,a) \text{ for all } x, a \Big\}.$$

Note that σ and $R_{\max}$ are allowed to change over contexts and actions. Formally, let an estimator be any function $\hat{v} : (\mathcal{X} \times \mathcal{A} \times \mathbb{R})^n \to \mathbb{R}$ that takes n data points collected by μ and outputs an estimate of v^π. The *minimax risk* of off-policy evaluation over the class $\mathcal{R}(\sigma^2, R_{\max})$ is defined as

$$R_n(\pi; \lambda, \mu, \sigma, R_{\max}) \coloneqq \inf_{\hat{v}} \sup_{D(r|x,a) \in \mathcal{R}(\sigma, R_{\max})} \mathbb{E}\left[(\hat{v} - v^\pi)^2 \right]. \tag{3}$$

Recall that the expectation here is taken over the n samples collected according to μ, along with any randomness in the estimator. The main goal of this section is to obtain a lower bound on the minimax risk. To state our bound, recall that $\rho(x,a) = \pi(a \mid x)/\mu(a \mid x)$ is an importance weight at (x,a). We make the following technical assumption on our problem instances, described by tuples of the form $(\pi, \lambda, \mu, \sigma, R_{\max})$:

Assumption 1. *There exists $\epsilon > 0$ such that $\mathbb{E}_\mu\left[(\rho R_{\max})^{2+\epsilon}\right]$ and $\mathbb{E}_\mu\left[(\rho\sigma)^{2+\epsilon}\right]$ are finite.*

This assumption is fairly mild, as it is only a slight strengthening of the assumption that $\mathbb{E}_\mu[(\rho R_{\max})^2]$ and $\mathbb{E}_\mu[(\rho\sigma)^2]$ be finite, which is required for consistency of IPS (see, e.g., the bound on the variance of IPS in Dudík et al. [2014], which assumes the finiteness of these second moments). Our assumption holds for instance whenever the context space is finite and π cannot pick actions that receive zero probability under μ,

since both ρ and $R_{\max}$ are bounded in this case. The latter assumption is quite standard and important in the off-policy evaluation literature as we observed in Section 2.1.

3.2 Minimax lower bound for off-policy evaluation

With the minimax setup in place, we now give our main lower bound on the minimax risk of off-policy evaluation and discuss its conclusions. In order to describe our result, we define one convenient piece of notation

$$C := 2^{2+\epsilon} \cdot \max\left\{ \frac{\mathbb{E}_\mu\left[(\rho R_{\max})^{2+\epsilon}\right]^2}{\mathbb{E}_\mu\left[(\rho R_{\max})^2\right]^{2+\epsilon}}, \frac{\mathbb{E}_\mu\left[(\rho\sigma)^{2+\epsilon}\right]^2}{\mathbb{E}_\mu\left[(\rho\sigma)^2\right]^{2+\epsilon}} \right\} \tag{4}$$

Theorem 1 (Minimax lower bound)**.** *For any problem instance satisfying Assumption 1 with ϵ and any $n \geq \max\left\{5C^{1/\epsilon},\, C^{2/\epsilon}\mathbb{E}_\mu[\sigma^2/R_{\max}^2]\right\}$, the minimax risk satisfies the lower bound*

$$R_n(\pi; \lambda, \mu, \sigma, R_{\max}) \geq \frac{1}{220n}\left(\mathbb{E}_\mu\left[\rho^2\sigma^2\right] + \mathbb{E}_\mu\left[\rho^2 R_{\max}^2\right]\left(1 - 110\lambda_0 \log(4/\lambda_0)\right)\right) \quad ,$$

where $\lambda_0 = \max_{x\in\mathcal{X}} \lambda(x)$ is the largest probability of a single context.

The theorem is proved in Appendix A and uses a reduction to hypothesis testing, followed by an argument due to Le Cam. We now discuss several aspects of the above result.

Pre-conditions of the theorem The theorem assumes the existence of a (problem-dependent) constant C which depends on variance moments of the importance weighted rewards. The stated lower bound holds as long as the number of samples n is large enough relative to C and some other parameters. Assuming that the constant C is finite, then the condition on n is eventually satisfied as long as the random variable $\sigma/R_{\max}$ has a bounded variance. We believe this is quite reasonable, since the problems of interest typically have heavy tails in the importance weights but not the reward functions. In terms of the importance weights, we require the random variables $\rho R_{\max}$ and $\rho\sigma$ to have appropriate moment bounds. In particular, whenever Assumption 1 is satisfied, then it is easy to see that $C < \infty$. In particular, under tail decay conditions on these two random variables, the precondition of the theorem is met for a reasonable number of samples. An extreme case where the preconditions will be violated are if $R_{\max} \equiv 0$, in which case the estimation problem is indeed trivial. For the remainder of this discussion, we will assume that n is appropriately large so that the preconditions of the theorem are met.

Comparison with upper bounds In typical contextual bandit settings, the contexts are drawn from either a continuous distribution, or an extremely large, but finite set. In either case, the probability of observing any particular context is vanishingly small, so that it is reasonable to expect $\lambda_0 \to 0$ as $n \to \infty$. Consequently, the lower bound becomes $\Omega\left(\mathbb{E}_\mu[\rho^2(\sigma^2 + R_{\max}^2)]/n\right)$. Comparing with the upper bounds on the MSE for existing estimators, we find that this lower bound precisely matches the upper bound for the IPS estimator [Horvitz and Thompson, 1952] (defined in Equation 2). We present the upper bound here for completeness.

Lemma 1 (IPS estimator [Horvitz and Thompson, 1952])**.** *With $\hat{v}^\pi_{\mathrm{IPS}}$ defined in Equation 2, we have*

$$\mathbb{E}(\hat{v}^\pi_{\mathrm{IPS}} - v^\pi)^2 \leq \frac{1}{n}\left(\mathbb{E}_\mu[\rho^2\sigma^2] + \frac{1}{4}\mathbb{E}_\mu[\rho^2 R_{\max}^2]\right).$$

We conclude that IPS is unimprovable, in the worst case, beyond constant factors. Another implication is that the lower bound is sharp for our setting, and cannot be improved any further unless we change the minimax framework to consider smaller classes of reward distributions or allow other side information. In particular, we can combine the upper and lower bounds for a precise characterization of the minimax risk.

Corollary 1. *Under conditions of Theorem 1, for sufficiently small λ_0 and large enough n:*

$$\inf_{\hat{v}} \sup_{D(r|a,x)\in\mathcal{R}(\sigma^2,R_{\max})} \mathbb{E}_\mu(\hat{v}-v^\pi)^2 = \Theta\left[\frac{1}{n}\left(\mathbb{E}_\mu[\rho^2\sigma^2]+\mathbb{E}_\mu[\rho^2 R_{\max}^2]\right)\right].$$

Comparison with the multi-armed bandit setting of Li et al. [2015] The most closely related prior result was due to Li et al. [2015], who showed matching upper and lower bounds on the minimax risk for multi-armed bandits. The somewhat surprising conclusion of their work was the sub-optimality of IPS, which might appear at odds with our conclusion regarding IPS above. However, this difference actually highlights the additional challenges in contextual bandits beyond multi-armed bandits. This is best illustrated in a noiseless setting, where $\sigma = 0$ in the rewards. This makes the multi-armed bandit problem trivial, we can just measure the reward of each arm with one pull and find out the optimal choice. However, there is still a non-trivial lower bound of $\Omega(\mathbb{E}_\mu[\rho^2 R_{\max}^2]/n)$ in the contextual bandit setting, which is exactly the upper bound on the MSE of IPS when the rewards have no noise.

This difference crucially relies on λ_0 being suitably small relative to the sample size n. When the number of contexts is small, independent estimation for each context can be done in a noiseless setting as observed by Li et al. [2015]. However, once the context distribution is rich enough, then even with noiseless rewards, there is significant variance in the value estimates based on which contexts were observed. This distinction is further highlighted in the proof of Theorem 1, and is obtained by combining two separate lower bounds. The first lower bound considers the case of noisy rewards, and is a relatively straightforward generalization of the proof of Li et al. [2015]. The second lower bound focuses on noiseless rewards, and shows how the variance in a rich context distribution allows the environment to essentially simulate noisy rewards, even when the reward signal itself is noiseless.

4 Incorporating model-based approaches in policy evaluation

Amongst our twin goals of optimal and adaptive estimators, the discussion so far has centered around the optimality of the IPS estimators in a minimax sense. However, real datasets seldom display worst-case behavior, and in this section we discuss approaches to leverage additional structure in the data, when such knowledge is available. We begin with the necessary setup, before introducing our new estimator and its properties. Throughout this section, we drop the superscript π from value estimators, as the evaluation policy π is fixed throughout this discussion.

4.1 The need for model-based approaches

As we have seen in the last section, the model-free approach has a information-theoretic limit that depends quadratically on $R_{\max}$, σ and importance weight ρ. This is good news in that it implies the existence of an optimal estimator–IPS. However, it also substantially limits what policies can be evaluated, due to the quadratic dependence on ρ. If $\mu(a \mid x)$ is small for some actions, as is typical when number of actions is large or in real systems with a cost for exploration, the policies π which can be reliably evaluated cannot put much mass on such actions either without resulting in unreliable value estimates. The key reason for this limitation is that the setup so far allows completely arbitrary reward models—$\mathbb{E}[r \mid x, a]$ can change arbitrarily across different actions and contexts. Real data sets are seldom so pathological, and often we have substantial intuition about contexts and actions which obtain similar rewards, based on the specific application. It is natural to ask how can we leverage such prior information, and develop better estimators that improve upon the minimax risk in such favorable scenarios.

In the presence of additional domain information, a common approach is to construct an explicit model for the expected reward, given context and action. Based on the existing logs according to μ: $(x_i, a_i, r_i)_{i=1}^n$, one can then fit the parameters in such a model to form a reward estimator $\hat{r}(x, a)$. The task of policy evaluation is now simply performed by scoring π according to $\hat{r}$ as

$$\hat{v}_{\text{DM}} = \frac{1}{n} \sum_{i=1}^{n} \sum_{a \in \mathcal{A}} \pi(a|x_i) \hat{r}(x_i, a), \tag{5}$$

where the DM stands for *direct method*, a name for this approach which has been used in prior works [Dudík et al., 2014]. This approach appears attractive when $\hat{r}$ is a close to $\mathbb{E}[r \mid x, a]$. Crucially, $\hat{r}$ can be evaluated for any x, a pair, meaning that there is no need to do use importance weights unlike in IPS. This suggests that we might completely eliminate any dependence on ρ using this approach., and it is easily seen that given any estimator $\hat{r}$ such th which takes values in $[0, R_{\max}(x, a)]$ for any x, a, the variance satisfies:

$$\text{Var}(\hat{v}_{\text{DM}}) \leq \frac{1}{n} \mathbb{E}_\pi[R_{\max}^2] = \frac{1}{n} \mathbb{E}_\mu[\rho R_{\max}^2]. \tag{6}$$

That is, there is a linear, rather than quadratic dependence on ρ, unlike in the worst case minimax risk of Corollary 1. The catch, however, is that such estimators can have an uncontrolled bias of $\mathcal{O}(\mathbb{E}_\pi[R_{\max}])$ in the worst case, meaning that even asymptotic consistency is not guaranteed.

Prior works have addressed the bias of DM, while trying to reduce the variance of IPS by using *doubly robust* estimators. A concrete estimator previously used for off-policy evaluation in Dudík et al. [2014] is defined as

$$\hat{v}_{\text{DR}} = \frac{1}{n} \sum_{i=1}^{n} \left[\frac{(r_i - \hat{r}(x_i, a_i))\pi(a_i)}{\mu(a_i)} + \sum_{a \in \mathcal{A}} \pi(a|x_i) \hat{r}(x_i, a_i) \right]. \tag{7}$$

The theory of these methods suggests that the variance of $\hat{v}_{\text{DR}}^{\pi}$ can be smaller than IPS, if $\hat{r}$ is a good reward estimator. Does this imply that DR is the right model-based adaptive estimator?

An interesting benchmark to consider is when one has a prefect reward estimator, that is $\hat{r}(x, a) = \mathbb{E}[r \mid x, a]$ almost surely. In this case, $\hat{v}_{\text{DM}}$ has no bias, and hence its MSE is identical to its variance, with a linear dependence on importance weights as in Equation 6. However, based on the results of Dudík et al. [2014], the MSE of the DR estimator in this special case is

$$\text{MSE}(\hat{v}_{\text{DR}}) \leq \frac{1}{n} \left(\mathbb{E}_\mu[\rho^2 \sigma^2] + \mathbb{E}_\mu[\rho R_{\max}^2] \right). \tag{8}$$

Crucially, the MSE still has a dependence on ρ^2 in it as long as the variance σ is non-trivial. Comparing Equations 6 and 8, it would be better to use $\hat{v}_{\text{DM}}$ instead in this case. Also note that ρ instead of ρ^2 term multiplying $R_{\max}$ does not contradict the lower bound in Theorem 1, since the bound (8) assumes access to the expected reward function which is information an estimator does not have in our minimax setup and which cannot be computed from the samples in general.

We conclude that while DR improves upon IPS in some cases, the estimator still does not adapt enough in the presence of good reward models, when the rewards have a non-trivial variance. We next present an estimator which utilizes the reward model in a different way, with an eye towards a more favorable bias-variance tradeoff.

4.2 The SWITCH estimators to incorporate reward models

We now present a class of estimators which incorporate any available reward models very differently from the DR approach. The starting point for our method is the observation that insistence on maintaining unbiasedness

puts the DR estimator at one extreme end of the bias-variance tradeoff. Prior works have considered ideas such as truncating the rewards, or importance weights when the importance weights get large (see e.g. Bottou et al. [2013] for a detailed discussion), which can often reduce the variance drastically at the cost of a little bias. We take the intuition a step further, and propose to estimate the rewards for actions differently, based on whether they have a large or a small importance weight given a context. When importance weights are small, we continue to use our favorite unbiased estimators, but switch to using the (potentially biased) reward model on the actions with large importance weights. Here small and large are defined via some *threshold parameter* τ. Varying this parameter between 0 and ∞ leads to a family of estimators which we call the SWITCH estimators as they switch between a model-free and model-based approach.

We now formalize this intuition, and begin by decomposing the value of π according to importance weights:

$$\begin{aligned} v^\pi &= \mathbb{E}_\pi[r] = \mathbb{E}_\pi[r\mathbf{1}(\rho \le \tau)] + \mathbb{E}_\pi[r\mathbf{1}(\rho > \tau)] \\ &= \mathbb{E}_\mu[\rho r\mathbf{1}(\rho \le \tau)] \;+\; \mathbb{E}_{x\sim\lambda}\Big[\sum_{a\in\mathcal{A}} \mathbb{E}_D[r|x,a]\,\pi(a|x)\,\mathbf{1}(\rho(x,a) > \tau)\Big]. \end{aligned}$$

Conceptually, we split our problem into two. The first problem always has small importance weights, so we can use unbiased estimators such as IPS or DR as before. The second problem, where importance weights are large, is essentially addressed by DM. Writing this out leads to the following estimator:

$$\hat{v}_{\text{SWITCH}} = \frac{1}{n}\sum_{i=1}^{n}[r_i\rho_i\mathbf{1}(\rho_i \le \tau)] + \frac{1}{n}\sum_{i=1}^{n}\sum_{a\in\mathcal{A}}\hat{r}(x_i,a)\pi(a|x_i)\mathbf{1}(\rho(x_i,a) > \tau). \tag{9}$$

Note that the above estimator specifically uses IPS on the first part of the problem. We will mention an alternative using the DR estimator for the first part at the end of this section. We first present a bound on the MSE of the SWITCH estimator using IPS, for a given choice of the threshold τ.

Theorem 2. *Let $\epsilon(a,x) := \hat{r}(a,x) - \mathbb{E}[r|a,x]$ be the bias of $\hat{r}$ and assume $\hat{r}(x,a) \in [0, R_{\max}(x,a)]$ almost surely. Then for every $n = 1,2,3,\ldots$, and for the $\tau > 0$ used in Equation 9, we have*

$$\text{MSE}(\hat{v}_{\text{SWITCH}}) \le \frac{2}{n}\left\{\mathbb{E}_\mu\big[\left(\sigma^2 + R_{\max}^2\right)\rho^2\mathbf{1}(\rho \le \tau)\big] + \mathbb{E}_\mu\left[\rho R_{\max}^2\mathbf{1}(\rho > \tau)\right] + \mathbb{E}_\mu\big[\rho\epsilon \;\big|\; \rho > \tau\big]^2\pi(\rho > \tau)^2\right\},$$

where quantities $R_{\max}$, σ, ρ, and ϵ are functions of the random variables x and a, and we recall the use of π and μ as joint distributions over (x,a,r) tuples.

Remark 1. The proposed estimator is an interpolation of the DM and IPS estimators. By taking $\tau =$, we get SWITCH coincides with DM while $\tau \to \infty$ yields IPS. Several estimators related to SWITCH have been studied in the literature, and we discuss a couple of them here.

- A special case of SWITCH uses $\hat{r} \equiv 0$, meaning that all the actions with large importance weights are essentially eliminated from consideration. This approach, with a specific choice of τ was described in Bottou et al. [2013] and will be evaluated in the experiments under the name *Trimmed IPS*.
- Thomas and Brunskill [2016] study a similar estimator in the more general context of reinforcement learning. Their approach can be seen as using a number of candidate threshold τ's and then evaluating the policy as a weighted sum of the estimates corresponding to each τ. They address the questions of picking these thresholds and the weights in a specific manner for their estimator called MAGIC, and we discuss these aspects in more detail in the following subsection.

Remark 2. The MSE bound in Theorem 2 is easily interpreted. The first term is the MSE of IPS on the region where it is used, while the second and third terms capture the variance and squared bias of DM on its region respective. Consequently, our matches the result for IPS when $\tau \to \infty$, meaning that the SWITCH estimators are also minimax optimal when τ is appropriately chosen. At the other extreme, we can consider the case of $\tau = 0$ and a perfect reward predictor so that $\epsilon \equiv 0$. In this case, the MSE bound matches that of DM in Equation 6. More generally, the estimators are robust to heavy-tails in the distribution of importance weights, unlike both IPS and DR estimators.

Remark 3. The policy value in the region where the importance weights are small can be estimated using any unbiased approach rather than just IPS. For instance, we can use the DR, giving rise to the estimator, which we denote SWITCH-DR. For any reward estimator $\hat{r}'$ to construct the DR estimator, we get

$$\begin{aligned}\hat{v}_{\text{SWITCH-DR}} =& \frac{1}{n}\sum_{i=1}^{n}\left[(r_i - \hat{r}'_i)\rho_i \mathbf{1}(\rho_i \le \tau) + \sum_{a\in\mathcal{A}} \hat{r}'(x_i,a)\pi(a|x_i)\mathbf{1}(\rho(x_i,a) \le \tau)\right] \\ &+ \frac{1}{n}\sum_{i=1}^{n}\sum_{a\in\mathcal{A}} \hat{r}(x_i,a)\pi(a|x_i)\mathbf{1}(\rho(x_i,a) > \tau).\end{aligned} \tag{10}$$

Note that $\hat{r}$ and $\hat{r}'$ need not be the same direct estimators. When they are indeed the same, $\hat{v}_{\text{SWITCH-DR}}$ can be rewritten as essentially a modified doubly robust estimator where all the importance weights above τ are clipped to zero, so that the importance weighted part of DR makes no contribution.

The analysis in Theorem 2 still applies, replacing the variance of IPS with that of DR from Dudík et al. [2014]. Since no independence was required in our analysis between the IPS and the DM parts of the estimator, the result is also robust to the use of a common data-dependent estimator $\hat{r} = \hat{r}'$ in SWITCH-DR (10).

4.3 Automatic parameter tuning

So far we have discussed the properties of the SWITCH estimators, assuming that the parameter τ is chosen well. On the other hand, the estimator can be as bad as DM in the worst-case if τ is poorly chosen, as evidenced by the bias term in Theorem 2. For the estimators to be useful in practice, it becomes essential to have a procedure to select a good value of τ. A natural criterion for selecting τ would be to pick one that minimizes the MSE of the resulting estimator. Since we do not know the precise MSE (as v^π is unknown), an alternative is to minimize a data-dependent estimate for it. Recalling that the MSE can be written as the sum of variance and squared bias, we estimate and bound the terms individually.

We can estimate the variance of the SWITCH estimator in a straightforward manner from the data. Let $Y_i(\tau)$ denote the estimated value that π obtains on the data point x_i according to the SWITCH estimator with the threshold τ, that is

$$Y_i(\tau) := r_i\rho_i\mathbf{1}(\rho_i \le \tau) + \sum_{a\in\mathcal{A}} \hat{r}(x_i,a)\pi(a|x_i)\mathbf{1}(\rho(x_i,a) > \tau) \quad \text{and} \quad \bar{Y}(\tau) = \frac{1}{n}\sum_{i=1}^{n} Y_i(\tau),$$

Since the x_i are i.i.d., the variance of $\hat{v}_{\text{SWITCH}}$ can be simply estimated as

$$\text{Var}(\hat{v}_{\text{SWITCH}-\tau}) = \frac{1}{n}\text{Var}(\hat{v}_{\text{SWITCH}-\tau}(x_1)) \approx \frac{1}{n^2}\sum_{i=1}^{n}(Y_i(\tau) - \bar{Y}(\tau))^2 =: \widehat{\text{Var}}_\tau, \tag{11}$$

where the approximation above is clearly consistent since the random variables Y_i are appropriately bounded as long as the rewards are bounded, with the importance weights capped at the threshold τ. Note that we have explicitly denoted the dependence on τ in the SWITCH estimator above.

Next we turn to the bias term. For understanding bias, we look at the MSE bound in Theorem 2, and observe that the last term in that theorem is precisely an upper bound on the bias. Rather than using a direct bias estimate, which would require knowledge of the error in $\hat{r}$, we will upper bound this term. For now, let us assume that the function $R_{\max}(x, a)$ is known. In many practical applications, this is not a limiting assumption at all, where an apriori bound on the rewards is known ahead of time anyways. Then we can upper bound the bias term as

$$\begin{aligned}\text{Bias}^2(\hat{v}_{\text{SWITCH}}) &\leq \mathbb{E}_\mu[\rho\epsilon^2|\rho > \tau]\pi(\rho > \tau)^2 \leq \mathbb{E}_\mu[\rho R_{\max}^2|\rho > \tau]\pi(\rho > \tau)^2 \\ &\approx \left[\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}_\pi\left(R_{\max}^2|\rho > \tau, x_i\right)\right]\left[\frac{1}{n}\sum_{i=1}^{n}\pi(\rho > \tau|x_i)\right]^2 =: \widehat{\text{Bias}}_\tau^2. \end{aligned} \tag{12}$$

Note that this estimate of the squared bias is not as straightforward as the one for variance, since we independently use sample-based unbiased approximations of the expectation of $R_{\max}$ and the probability of large importance weights, and then multiply them. However, standard arguments can be used to show that the resulting estimates are still consistent under mild conditions. As a special case, if $R_{\max} \equiv R$, where the uniform upper bound R is known, then the first average simplifies to R making the bias estimate rather straightforward.

With these estimates, we pick the value of τ that minimizes

$$\widehat{\tau} = \operatorname*{argmin}_\tau \widehat{\text{Var}}_\tau + \widehat{\text{Bias}}_\tau^2. \tag{13}$$

Notice that the upper bound on the bias is rather conservative in that it upper bounds the error of DM at the largest possible value for every data point. This has the effect of favoring the use of the unbiased part in SWITCH whenever possible, unless the variance would overwhelm even an arbitrarily biased DM. While this might appear overly conservative, this does imply the minimax optimality of the SWITCH estimator using $\widehat{\tau}$ almost immediately: the incurred bias is no more than our upper bound, and we incur it only when the minimax optimal IPS estimator would be suffering an even larger variance.

Finally, we should mention that this development is quite related to the MAGIC estimator of Thomas and Brunskill [2016], which was discussed following Theorem 2. The key differences are that we pick only one threshold τ, while they compute the estimates with many different τ's. Rather than picking just one of these estimates, they further learn a more general weighting over the estimates. In that sense, our estimator can be seen as a special case which puts the entire weight on one choice. Like us, they also pick this weighting function by optimizing a bias variance tradeoff. However, we use very different estimates for bias and variance than their estimator. In the experiments, we did evaluate their approach for picking the threshold τ and found that the choice $\widehat{\tau}$ in Equation 13 generally works better.

5 Experiments

In this section, we conduct an experimental evaluation of the proposed SWITCH estimators. We will be using the same 10 UCI data sets that was used previously by Dudík et al. [2011] and convert the multi-class classification problem to contextual bandits by

1. making it a sequential learning task, where the learner receives one example at a time, and guesses its label using policy μ.

2. providing a reward of 1 (and hence revealing the correct label) only if the guess is correct; otherwise providing a 0-reward. Note that this means $R_{\max} \equiv 1$ is a valid bound.

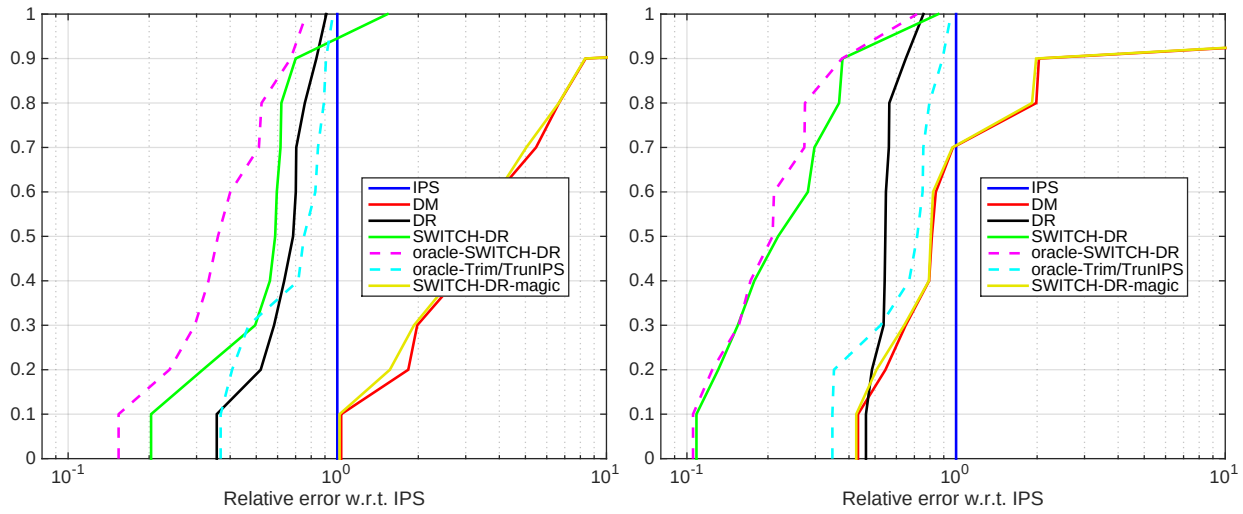

(a) Deterministic reward (b) Noisy reward

Figure 1: Cumulative distribution of the Rel. MSE of various estimators over 10 UCI data sets. In the left-hand panel, the UCI labels are used as is; in the right-hand panel, additional label noise is added. Curves towards top-left part of the plot are better, as they achieve smaller values of MSE. Methods in dashed lines are "cheating" by choosing the threshold parameter τ to optimize test MSE. As we see, SWITCH-DR dominates all other methods and our empirical tuning of τ is not too far from the optimal possible. Each dataset uses an n which is the actual size of the dataset, drawn via bootstrap sampling and results are averaged over 500 trials.

The target policy π we used is the deterministic decision of a learned logistic regression classifier, while the logging policy μ we used is the probability estimates of a learned logistic regression classifier on a covariate-shifted data set. We simulate the covariate-shifted data set using the same technique as in Dudík et al. [2011], which follows standard practice as was described by Gretton et al. [2009].

In each data set with n rows, we treat the uniform distribution over the data set itself as a surrogate of the population distribution so that we know the ground truth of the rewards. Then, in the simulator, we randomly draw i.i.d. data sets of size $[100, 200, 500, 1000, 2000, 5000, 10000, ...]$ until it reaches n, each repeating for 500 times, then compare different candidate estimators' MSE, estimated using the empirical averages of the square error over the 500 replicates, which we can calculate exactly since we know v^π. It is worth pointing out that for some of the baselines that we compare to, e.g., IPS and DR, their distribution of square errors can have very large variance due to the potentially large importance weights. This leads to very large error bars if we aggregate their MSE with even 500 replicates. To circumvent this issue, we report a clipped version of the MSE that truncates any square errors to 1 if it is larger than 1, namely

$$\text{MSE} = \mathbb{E}[(\hat{v} - v^\pi)^2 \wedge 1].$$

This allows us to get valid confidence intervals for our empirical estimates of this quantity. Note that this does not change the MSE estimate of our approach at all, but is significantly more favorable towards IPS and DR. In this section, whenever we refer to"MSE", we are referring to the truncated version.

We will compare the SWITCH-IPS and SWITCH-DR against the following baselines:

1. IPS estimator
2. Direct method via logistic regression

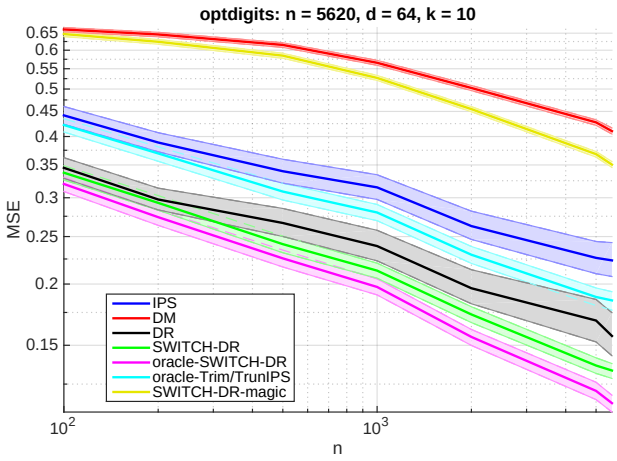

(a) optdigits / deterministic reward

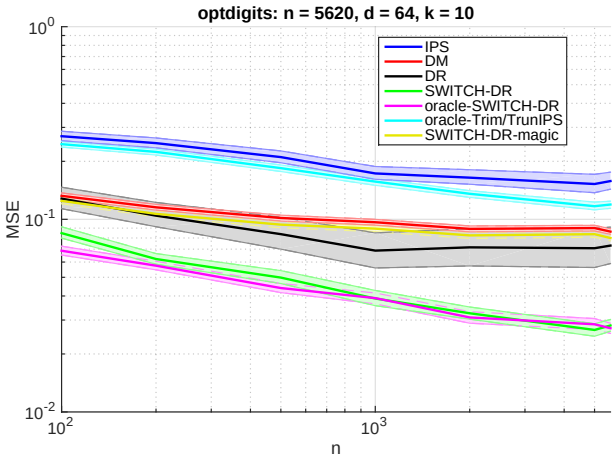

(b) optdigits / noisy reward

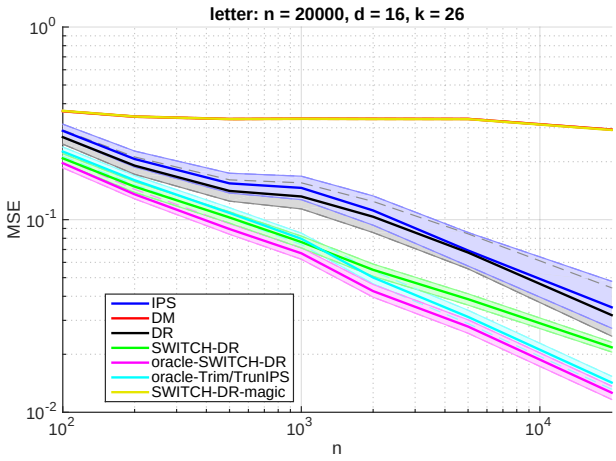

(c) letter / deterministic reward

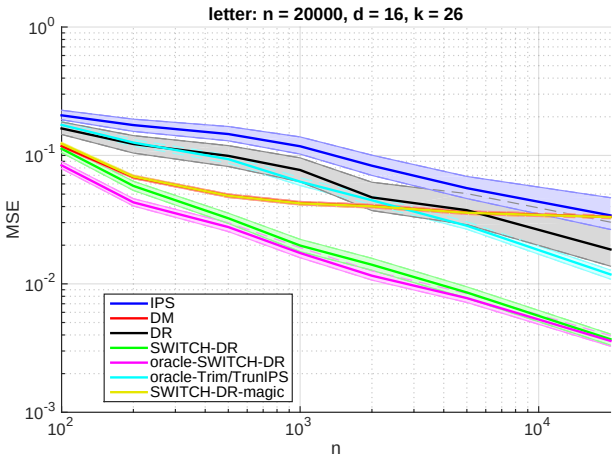

(d) letter / noisy reward

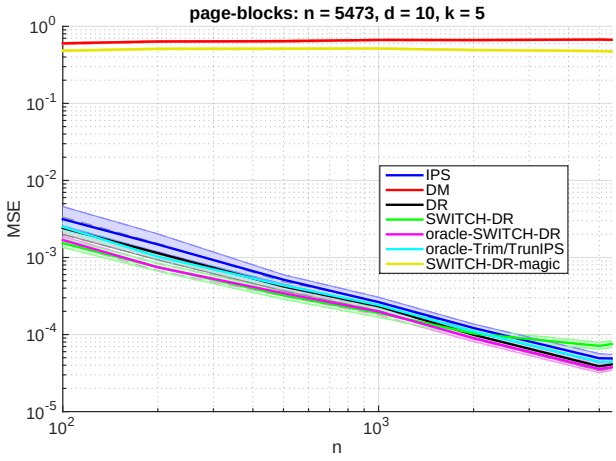

(e) page-blocks / deterministic reward

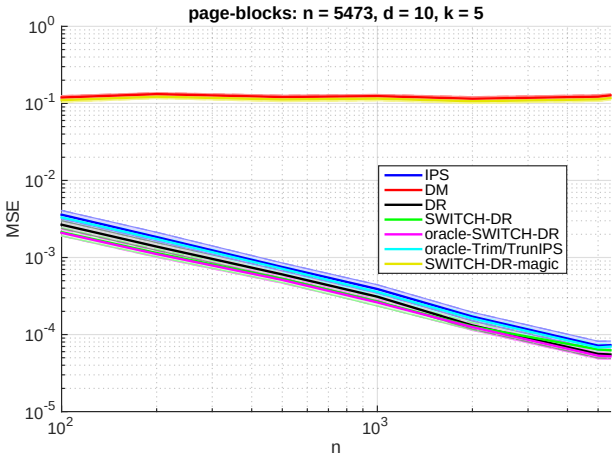

(f) page-blocks / noisy reward

Figure 2: MSE of different methods over input data size. The first row is the "optdigit" dataset, the second row is the "letter" data set, and the third row are the results from the "page-blocks" data set.

3. Doubly robust estimator
4. Truncated and reweighted IPS
5. Trimmed IPS

The doubly robust is constructed by randomly splitting the data set into two halves, estimating DM on one half, combining with IPS on the other half, then switching the two sets and doing it again. The final estimate is the average of the two. This approach was suggested by Dudík et al. [2011], and is standard practice of doubly robust estimators when the direct method (or the oracle estimator) needs to be estimated from the data.

We also consider two additional variants of IPS, where importance weights are either capped at τ and renormalized [see, e.g., Bembom and van der Laan, 2008], or the terms with weights larger than τ are removed altogether as described in Bottou et al. [2013]. Note that the trimmed IPS is a special case of SWITCH when the direct estimator $\equiv 0$.

For both our methods (as well as the two IPS variants) we tune the parameter τ by the method described in Section 4.3. So everything is implementable in practice. As a reference, we are also including in the plot the results for the optimally tuned τ from hindsight, as well as the alternative ensemble approach via MAGIC. This allows us to gauge how well our automatic parameter tuning procedure works on these data sets.

In addition, we repeated the same experiments for a more challenging setting where the feedbacks are noisy. This is simulated by only revealing the correct reward with probability 0.5 and outputting a random coin toss with probability the other 0.5. Theoretically, this should lead to bigger σ^2 and larger variance in all estimators.

In order to stay comparable across data sets, we use relative MSE with respect to the IPS. For each estimator $\hat{v}$, report the relative MSE with respect to the IPS estimator

$$\text{Rel.MSE}(\hat{v}) = \frac{\text{MSE}(\hat{v})}{\text{MSE}(\hat{v}_{\text{IPS}})},$$

such that the numerical value is comparable across data sets and is independent to n.

The results are summarized in Figure 1[2], in which methods that achieve smaller values of MSE are towards the top-left corner of the plot. In general, we want to maximize the area under the curve as these are CDF plots, and we notice the logarithmic scale for the relative MSE values on the X-axis. As we see, SWITCH-DR dominates all other methods and our empirical tuning of τ is not too far from the optimal possible. The advantage of SWITCH-DR is even stronger in the noisy-reward setting, where we add label noise to UCI data.

In Figure 2, we illustrate the convergence of MSE as n increases. In particular, we selected three data sets and show how SWITCH and SWITCH-DR perform against baselines in three typical cases: (i) when the direct method works better than IPS and DR; (ii) when the direct method works well initially but then as n gets large outperformed by IPS and DR; (iii) when the direct method works extremely poorly. In the first two cases, SWITCH-DR outperforms both DM and IPS, while DR only improves over IPS slightly. In the third case, SWITCH-DR performs about as well as IPS and DR despite that DM is very bad. In all cases, SWITCH-DR is robust to additional noise in the reward.

6 Conclusion

In this paper we carried out minimax analysis of off-policy evaluation in contextual bandits and showed that IPS is optimal in the worst-case. This result highlights the need for using side information, potentially

[2]For the clarity of the presentation, we excluded SWITCH, which also significantly outperforms either IPS but is dominated by SWITCH-DR. For the same reason, we combined the Trim and Truncated/reweighted IPS.

provided by modeling the reward directly, especially when importance weights are too large. Given this observation, we proposed a new class of estimators called SWITCH that can be used to combine any importance sampling estimators, including IPS and DR, with DM. The estimator involves adaptively switching to DM when the importance weights are large and switching to either IPS or DR when the importance weights are small. We showed that the new estimator has favorable theoretical properties and also works well on real-world data.

Acknowledgments

The work was partially completed during YW's internship at Microsoft Research NYC from May 2016 - Aug 2016. The authors would like to thank Lihong Li and John Langford for helpful discussions, Edward Kennedy for bringing our attention to related problems and recent development in causal inference and Wei He for pointing us to [Sun, 2006] for rigorous a measure-theoretic treatment to a continuum of random variables and the corresponding law of large numbers.

YW was supported by NSF Award BCS-0941518 to CMU Statistics, a grant by Singapore NRF under its International Research Centre @ Singapore Funding Initiative, and a Baidu Scholarship.

A Proof of Theorem 1

In this appendix we prove the minimax bound of Theorem 1. The result is obtained by combining the following two lower bounds:

Theorem 3 (Lower bound 1)**.** *For each problem instance such that $\mathbb{E}_\mu[\rho^2\sigma^2] < \infty$, we have*

$$R_n(\pi;\lambda,\mu,\sigma,R_{\max}) \geq \frac{\log 2\mathbb{E}_\mu[\rho^2\sigma^2]}{8n}\left[1 - \frac{\mathbb{E}_\mu\left[\rho^2\sigma^2\mathbf{1}\left(\rho\sigma^2 > R_{\max}\sqrt{n\mathbb{E}_\mu[\rho^2\sigma^2]/(2\log 2)}\right)\right]}{\mathbb{E}_\mu[\rho^2\sigma^2]}\right]^2 .$$

Theorem 4 (Lower bound 2)**.** *For each problem instance such that $\mathbb{E}_\mu[\rho^2 R_{\max}^2] < \infty$, we have*

$$\begin{aligned} &R_n(\pi;\lambda,\mu,\sigma,R_{\max}) \\ &\geq \frac{\log 2\mathbb{E}_\mu[\rho^2 R_{\max}^2]}{8n}\left[1 - \frac{8\lambda_0\log(4/\lambda_0)}{\log 2} - \frac{\mathbb{E}_\mu\left[\rho^2 R_{\max}^2\mathbf{1}\left(\rho R_{\max} > \sqrt{n\mathbb{E}_\mu[\rho^2 R_{\max}^2]/(16\log 2)}\right)\right]}{\mathbb{E}_\mu[\rho^2 R_{\max}^2]}\right]^2 , \end{aligned}$$

where $\lambda_0 = \max_{x\in\mathcal{X}}\lambda(x)$ is the largest probability of a single context.

The first bound captures the intrinsic difficulty in the variance of reward. The second result shows the additional dependence on $R_{\max}^2$, even when $\sigma \equiv 0$, whenever the distribution λ is sufficiently spread out over contexts, as quantified by the largest context probability λ_0. We next show how these two lower bounds yield Theorem 1 and then return to their proofs.

Proof of Theorem 1. We begin by simplifying the two lower bounds. Assume that Assumption 1 holds with ϵ and let $p = 1 + \epsilon/2$ and $q = 1 + 2/\epsilon$, i.e., $1/p + 1/q = 1$. Then the definition of C means that

$$\begin{aligned} C^{1/(\epsilon q)} = C^{1/(2+\epsilon)} &= 2\cdot\max\left\{\frac{\mathbb{E}_\mu\left[(\rho^2 R_{\max}^2)^{\frac{2+\epsilon}{2}}\right]^{\frac{2}{2+\epsilon}}}{\mathbb{E}_\mu\left[\rho^2 R_{\max}^2\right]}, \frac{\mathbb{E}_\mu\left[(\rho^2\sigma^2)^{\frac{2+\epsilon}{2}}\right]^{\frac{2}{2+\epsilon}}}{\mathbb{E}_\mu\left[\rho^2\sigma^2\right]}\right\} \\ &= 2\cdot\max\left\{\frac{\mathbb{E}_\mu\left[(\rho^2 R_{\max}^2)^p\right]^{1/p}}{\mathbb{E}_\mu\left[\rho^2 R_{\max}^2\right]}, \frac{\mathbb{E}_\mu\left[(\rho^2\sigma^2)^p\right]^{1/p}}{\mathbb{E}_\mu\left[\rho^2\sigma^2\right]}\right\} , \end{aligned} \tag{14}$$

and the assumption on n implies that

$$n \geq \max\left\{(16\log 2)C^{1/\epsilon},\ (2\log 2)C^{2/\epsilon}\mathbb{E}_\mu[\sigma^2/R_{\max}^2]\right\} . \tag{15}$$

First, we simplify the correction term in the lower bound of Theorem 3. Using Hölder's inequality and Eq. (14), we have

$$\begin{aligned} &\mathbb{E}_\mu\left[\rho^2\sigma^2\mathbf{1}\left(\rho\sigma^2 > R_{\max}\sqrt{n\mathbb{E}_\mu[\rho^2\sigma^2]/(2\log 2)}\right)\right] \\ &\leq \mathbb{E}_\mu\left[(\rho^2\sigma^2)^p\right]^{1/p}\cdot\mathbb{P}_\mu\left[\rho\sigma^2 > R_{\max}\sqrt{n\mathbb{E}_\mu[\rho^2\sigma^2]/(2\log 2)}\right]^{1/q} \\ &\leq \frac{1}{2}\mathbb{E}_\mu[\rho^2\sigma^2]\cdot C^{1/(\epsilon q)}\cdot\mathbb{P}_\mu\left[\rho\sigma^2/R_{\max} > \sqrt{n\mathbb{E}_\mu[\rho^2\sigma^2]/(2\log 2)}\right]^{1/q} \end{aligned}$$

and by Markov's inequality, Cauchy-Schwartz inequality, and Eq. (15)

$$
\begin{aligned}
&\leq \frac{1}{2}\mathbb{E}_\mu[\rho^2\sigma^2] \cdot C^{1/(\epsilon q)} \cdot \left(\frac{\mathbb{E}_\mu\left[\rho\sigma \cdot (\sigma/R_{\max})\right]}{\sqrt{n\mathbb{E}_\mu[\rho^2\sigma^2]/(2\log 2)}}\right)^{1/q} \\
&\leq \frac{1}{2}\mathbb{E}_\mu[\rho^2\sigma^2] \cdot C^{1/(\epsilon q)} \cdot \left(\frac{\sqrt{\mathbb{E}_\mu[\rho^2\sigma^2]} \cdot \sqrt{\mathbb{E}_\mu[\sigma^2/R_{\max}^2]}}{\sqrt{n\mathbb{E}_\mu[\rho^2\sigma^2]/(2\log 2)}}\right)^{1/q} \\
&= \frac{1}{2}\mathbb{E}_\mu[\rho^2\sigma^2] \cdot \left(C^{2/\epsilon} \cdot \frac{2\log 2}{n} \cdot \mathbb{E}_\mu[\sigma^2/R_{\max}^2]\right)^{1/2q} \leq \frac{1}{2}\mathbb{E}_\mu[\rho^2\sigma^2] \ . \qquad (16)
\end{aligned}
$$

For the correction term in Theorem 4, we similarly have

$$
\begin{aligned}
&\mathbb{E}_\mu\left[\rho^2 R_{\max}^2 \mathbf{1}\left(\rho R_{\max} > \sqrt{n\mathbb{E}_\mu[\rho^2 R_{\max}^2]/(16\log 2)}\right)\right] \\
&\leq \mathbb{E}_\mu\left[\left(\rho^2 R_{\max}^2\right)^p\right]^{1/p} \cdot \mathbb{P}_\mu\left[\rho R_{\max} > \sqrt{n\mathbb{E}_\mu[\rho^2 R_{\max}^2]/(16\log 2)}\right]^{1/q} \\
&\leq \frac{1}{2}\mathbb{E}_\mu[\rho^2 R_{\max}^2] \cdot C^{1/(\epsilon q)} \cdot \mathbb{P}_\mu\left[\rho^2 R_{\max}^2 > n\mathbb{E}_\mu[\rho^2 R_{\max}^2]/(16\log 2)\right]^{1/q}
\end{aligned}
$$

and by Markov's inequality and Eq. (15)

$$
\begin{aligned}
&\leq \frac{1}{2}\mathbb{E}_\mu[\rho^2 R_{\max}^2] \cdot C^{1/(\epsilon q)} \cdot \left(\frac{\mathbb{E}_\mu[\rho^2 R_{\max}^2]}{n\mathbb{E}_\mu[\rho^2 R_{\max}^2]/(16\log 2)}\right)^{1/q} \\
&= \frac{1}{2}\mathbb{E}_\mu[\rho^2 R_{\max}^2] \cdot \left(C^{1/\epsilon} \cdot \frac{16\log 2}{n}\right)^{1/q} \leq \frac{1}{2}\mathbb{E}_\mu[\rho^2 R_{\max}^2] \ . \qquad (17)
\end{aligned}
$$

Using Eq. (16), the bound of Theorem 3 simplifies as

$$
\begin{aligned}
&R_n(\pi; \lambda, \mu, \sigma, R_{\max}) \\
&\geq \frac{\log 2\mathbb{E}_\mu[\rho^2\sigma^2]}{8n}\left[1 - \frac{\mathbb{E}_\mu\left[\rho^2\sigma^2 \mathbf{1}\left(\rho\sigma^2 > R_{\max}\sqrt{n\mathbb{E}_\mu[\rho^2\sigma^2]/(2\log 2)}\right)\right]}{\mathbb{E}_\mu[\rho^2\sigma^2]}\right]^2 \\
&\geq \frac{\log 2\mathbb{E}_\mu[\rho^2\sigma^2]}{8n}\left(1 - \frac{1}{2}\right)^2 = \frac{\log 2\mathbb{E}_\mu[\rho^2\sigma^2]}{32n} \ .
\end{aligned}
$$

Similarly, by Eq. (17), Theorem 4 simplifies as

$$
\begin{aligned}
&R_n(\pi; \lambda, \mu, \sigma, R_{\max}) \\
&\geq \frac{\log 2\mathbb{E}_\mu[\rho^2 R_{\max}^2]}{8n}\left[1 - \frac{8\lambda_0\log(4/\lambda_0)}{\log 2} - \frac{\mathbb{E}_\mu\left[\rho^2 R_{\max}^2 \mathbf{1}\left(\rho R_{\max} > \sqrt{n\mathbb{E}_\mu[\rho^2 R_{\max}^2]/(16\log 2)}\right)\right]}{\mathbb{E}_\mu[\rho^2 R_{\max}^2]}\right]^2 \\
&\geq \frac{\log 2\mathbb{E}_\mu[\rho^2 R_{\max}^2]}{8n}\left[1 - \frac{8\lambda_0\log(4/\lambda_0)}{\log 2} - \frac{1}{2}\right]^2 \\
&\geq \frac{\log 2\mathbb{E}_\mu[\rho^2 R_{\max}^2]}{8n}\left(\left(\frac{1}{2}\right)^2 - 2\cdot\frac{1}{2}\cdot\frac{8\lambda_0\log(4/\lambda_0)}{\log 2}\right) = \frac{\log 2\mathbb{E}_\mu[\rho^2 R_{\max}^2]}{32n}\left(1 - \frac{32\lambda_0\log(4/\lambda_0)}{\log 2}\right) \ ,
\end{aligned}
$$

where in the last line we used $(a-b)^2 \geq a^2 - 2ab$. Combining the two simplified bounds yields

$$\begin{aligned}
&R_n(\pi; \lambda, \mu, \sigma, R_{\max}) \\
&\quad \geq \frac{1}{2} \cdot \frac{\log 2\mathbb{E}_\mu[\rho^2\sigma^2]}{32n} + \frac{1}{2} \cdot \frac{\log 2\mathbb{E}_\mu[\rho^2 R_{\max}^2]}{32n} \left(1 - \frac{32\lambda_0 \log(4/\lambda_0)}{\log 2}\right) \\
&\quad \geq \frac{\log 2}{64n} \left[\mathbb{E}_\mu[\rho^2\sigma^2] + \mathbb{E}_\mu[\rho^2 R_{\max}^2] \left(1 - \frac{32}{\log 2} \cdot \lambda_0 \log(4/\lambda_0)\right)\right] \\
&\quad \geq \frac{1}{220n} \left[\mathbb{E}_\mu[\rho^2\sigma^2] + \mathbb{E}_\mu[\rho^2 R_{\max}^2] \Big(1 - 110\lambda_0 \log(4/\lambda_0)\Big)\right] . \qquad \square
\end{aligned}$$

It remains to prove Theorems 3 and 4. They are both proved by a reduction to hypothesis testing, and invoke Le Cam's argument to lower-bound the error in this testing problem. As in most arguments of this nature, the key novelty lies in the construction of an appropriate testing problem that leads to the desired lower bounds. Before proving the theorems, we recall the basic result of Le Cam which underlies our proofs. We point the reader to the excellent exposition of Lafferty et al. [2008, Section 36.4] on more details about Le Cam's argument.

Theorem 5 (Le Cam's method Lafferty et al., 2008, Theorem 36.8)**.** *Let $\mathcal{P}$ be a set of distributions. For any pair $P_0, P_1 \in \mathcal{P}$,*

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{4} \int [p_0^n \wedge p_1^n] dx \geq \frac{\Delta}{8} e^{-nD_{\mathrm{KL}}(P_0 \| P_1)}. \tag{18}$$

where $\Delta = d(\theta(P_0), \theta(P_1))$

While the proofs of the two theorems share a lot of similarities, they have to use reductions to slightly different testing problems given the different mean and variance constraints in the two results. We begin with the proof of Theorem 3, which has a simpler construction.

A.1 Proof of Theorem 3

The basic idea of this proof is to reduce the problem of policy evaluation to that of Gaussian mean estimation where there is a mean associated with each x, a pair. We now describe our construction.

Creating a family of problems Since we aim to show a lower bound on the hardness of policy evaluation in general, it suffices to show a particular family of hard problem instances, such that every estimator requires the stated number of samples on at least one of the problems in this family. Recall that our minimax setup assumes that π, μ and λ are fixed and the only aspect of the problem which we can design is the conditional reward distribution $D(r \mid x, a)$. For Theorem 3, this choice is further constrained to satisfy $\mathbb{E}[r \mid x, a] \leq R_{\max}(x, a)$ and $\mathrm{Var}(r \mid x, a) \leq \sigma^2$. In order to describe our construction, it will be convenient to define the shorthand $\mathbb{E}[r \mid x, a] = \eta(x, a)$. We will identify a problem in our family with the function $\eta(x, a)$ as that will be the only changing element in our problems. For a chosen η, the policy evaluation question boils down to estimating $v_\eta^\pi = \mathbb{E}[r(x, a)]$, where the contexts x are chosen according to λ, actions are drawn from $\pi(x, a)$ and the reward distribution $D_\eta(r \mid x, a)$ is a normal distribution with mean $\eta(x, a)$

$$D_\eta(r \mid x, a) = \mathcal{N}(\eta(x, a), \sigma^2).$$

Clearly this choice meets the variance constraint by construction, and satisfies the upper bound so long as $\eta(x, a) \leq R_{\max}(x, a)$ almost surely. Since the evaluation policy π is fixed throughout, we will drop the superscript and use v_η to denote $v_e r^\pi$ in the remainder of the proofs. With some abuse of notation, we also

use $\mathbb{E}_\eta[\cdot]$ to denote expectations where contexts and actions are drawn based on the fixed choices λ and μ corresponding to our data generating distribution, and the rewards drawn from η. We further use P_η to denote this entire joint distribution over (x, a, r) triples.

Given this family of problem instances, it is easy to see that for any pair of η_1, η_2 which are both pointwise upper bounded by $R_{\max}$, we have the lower bound:

$$R_n(\lambda, \pi, \mu, \sigma^2, R_{\max}) \geq \inf_{\hat{v}} \max_{\eta \in \eta_1, \eta_2} \mathbb{E}_\eta \Big[\underbrace{(\hat{v} - v_\eta)^2}_{\ell_\eta(\hat{v})} \Big],$$

where we have introduced the shorthand $\ell_{\eta_i}(\hat{v})$ to denote the squared error of $\hat{v}$ to v_{η_i}. For a parameter $\epsilon > 0$ to be chosen later, we can further lower bound this risk as

$$\begin{aligned} R_n(\lambda, \pi, \mu, \sigma^2, R_{\max}) &\geq \max_{\eta_i \in \eta_1, \eta_2} \mathbb{E}_\eta[\ell_\eta(\hat{v})] \geq \max_{\eta_i \in \eta_1, \eta_2} \epsilon \mathbb{P}_\eta(\ell_\eta \geq \epsilon) \\ &\geq \frac{\epsilon}{2} \left(\mathbb{P}_{\eta_1}(\ell_{\eta_1}(\hat{v}) \geq \epsilon) + \mathbb{P}_{\eta_2}(\ell_{\eta_2}(\hat{v}) \geq \epsilon) \right), \end{aligned} \tag{19}$$

where the last inequality lower bounds the maximum by the average. So far we have been working with an estimation problem. We next describe how to reduce this to a hypothesis testing problem.

Reduction to hypothesis testing For turning our estimation problem into a testing problem, the idea is to identify a pair η_1, η_2 such that they are far enough from each other so that any estimator which gets a small estimation loss can essentially identify whether the data generating distribution corresponds to P_{η_1} or P_{η_2}. In order to do this, we take any estimator $\hat{v}$ and identify a corresponding test statistic which maps $\hat{v}$ into one of η_1, η_2. The way to do this is essentially identified in Equation 19, and we describe it next.

Note that since we are constructing a hypothesis test for the specific family of distributions P_{η_1} and P_{η_2}, it is reasonable to consider test statistics which have knowledge of η_1 and η_2, and hence the corresponding distributions. Consequently, these tests also know the true policy values v_{η_1} and v_{η_2} and the only uncertainty is which of them gave rise to the observed data samples. Therefore, for any estimator $\hat{v}$, we can a associate a statistic $\phi(\hat{v}) = \operatorname{argmin}_\eta \{\ell_{\eta_1}(\hat{v}), \ell_{\eta_2}(\hat{v})\}$.

Given this hypothesis test, we are interested in its error rate $\mathbb{P}_\eta(\phi(\hat{v}) \neq \eta)$. We first relate the estimation error of $\hat{v}$ to the error rate of the test. Suppose for now that

$$\ell_{\eta_1}(\hat{v}) + \ell_{\eta_1}(\hat{v}) \geq 2\epsilon, \tag{20}$$

so that at least one of the losses is at least ϵ. Suppose further that $\phi(\hat{v}) = \eta_1$, so that we know $\ell_{\eta_1}(\hat{v}) \leq \ell_{\eta_2}(\hat{v})$. Then we are assured that $\ell_{\eta_1}(\hat{v}) \leq \epsilon$ assures the test's correctness by Equation 20, since the other loss is guaranteed to be at least ϵ in this condition. The reasoning for the loss under η_2 being small is similar, so that a valid upper bound on our test's error (under a uniform choice of η_1, η_2) is given by

$$\begin{aligned} \max_{\eta \in \eta_1, \eta_2} \mathbb{P}_\eta(\phi(\hat{v}) \neq \eta) &= \frac{1}{2} \left(\mathbb{P}_{\eta_1}(\phi(\hat{v}) \neq \eta_1) + \mathbb{P}_{\eta_2}(\phi(\hat{v}) \neq \eta_2) \right) \\ &\leq \frac{1}{2} \left(\mathbb{P}_{\eta_1}(\ell_{\eta_1}(\hat{v}) \geq \epsilon) + P_{\eta_2}(\ell_{\eta_2}(\hat{v}) \geq \epsilon) \right) \\ &\leq \frac{1}{\epsilon} R_n(\lambda, \pi, \mu, \sigma^2, R_{\max}), \end{aligned} \tag{21}$$

where the final inequality uses our earlier lower bound in Equation 19.

It remains to establish our earlier supposition (20) to finish the relating our testing and estimation problems. Assume for now that η_1 and η_2 are chosen such that

$$(v_{\eta_1} - v_{\eta_2})^2 \geq 4\epsilon. \tag{22}$$

Then an application of Cauchy-Shwartz inequality $(a+b)^2 \leq 2a^2 + 2b^2$ yields

$$4\epsilon \leq (v_{\eta_1} - v_{\eta_2})^2 \leq 2(\hat{v} - v_{\eta_1})^2 + 2(\hat{v} - v_{\eta_2})^2 \leq 2\ell_{\eta_1}(\hat{v}) + 2\ell_{\eta_2}(\hat{v}),$$

which yields the posited bound 19.

Invoking Le Cam's argument So far we have identified a hypothesis testing problem and a test statistic whose error is upper bounded in terms of the minimax risk of our problem. In order to complete the proof, we now place a lower bound on the error of this test statistic. Recall the result of Le Cam 18, which places an upper bound on the attainable error in any testing problem. In our setting, this translates to

$$\max_{\eta \in \eta_1, \eta_2} \mathbb{P}_\eta(\phi(\hat{v}) \neq \eta) \geq \exp\left(-nD_{KL}(P_{\eta_1} \,||\, P_{\eta_2})\right).$$

Since the distribution of the rewards is a spherical Gaussian, the KL-divergence is given by the squared distance between the means, scaled by the variance, that is

$$D_{KL}(P_{\eta_1} \,||\, P_{\eta_2}) = \mathbb{E}\left[\frac{(\eta_1(x,a) - \eta_2(x,a))^2}{2\sigma^2}\right],$$

where we recall that the contexts and actions are drawn from λ and μ respectively. Since we would like the probability of error in the test to be a constant, it suffices to choose η_1 and η_2 such that

$$\mathbb{E}\left[\frac{(\eta_1(x,a) - \eta_2(x,a))^2}{2\sigma^2}\right] \leq \frac{\log 2}{n}. \tag{23}$$

Picking the parameters So far, we have not made any concrete choices for η_1 and η_2, apart from some constraints which we have introduced along the way. Note that we have the constraints (22) and (23) which try to ensure that η_1 and η_2 are not too close that an estimator does not have to identify the true parameter, or too far that the testing problem becomes trivial. Additionally, we have the upper and lower bounds of 0 and $R_{\max}$ on η_1 and η_2. In order to reason about these constraints, it is convenient to set $\eta_2 \equiv 0$, and pick $\eta_1(x,a) = \eta_1(x,a) - \eta_2(x,a) = \Delta(x,a)$. We now write all our constraints in terms of Δ.

Note that v_{η_2} is now 0, so that the first constraint (22) is equivalent to

$$v_{\eta_1} = \mathbb{E}_{\eta_1}[\rho(x,a)r(x,a)] = \mathbb{E}_{\Delta}[\rho(x,a)r(x,a)] \geq 2\sqrt{\epsilon},$$

where the importance weighting function ρ is introduced since P_{η_1} is based on choosing actions according to μ and we seek to evaluate π. The second constraint (23) is also straightforward

$$\mathbb{E}\left[\frac{\Delta^2}{2\sigma^2}\right] \leq \frac{\log 2}{n}.$$

The bound constraints are respected by imposing $0 \leq \Delta(x,a) \leq R_{\max}(x,a)$ almost surely.

The minimax lower bound is then yielded by the largest ϵ in the constraint (22) such that the other two constraints can be satisfied. This gives rise to the following variational characterization of the minimax lower bound

$$\begin{aligned} \max_{\Delta} \quad & \epsilon \\ \text{such that} \quad & \mathbb{E}_{\Delta}[\rho(x,a)r(x,a)] \geq 2\sqrt{\epsilon}, \\ & \mathbb{E}\left[\frac{\Delta^2}{2\sigma^2}\right] \leq \frac{\log 2}{n} \quad \text{and} \\ & 0 \leq \Delta(x,a) \leq R_{\max}(x,a). \end{aligned}$$

Instead of finding the optimal solution, we show a feasible setting of Δ here. We set

$$\Delta = \min\left\{\frac{\alpha\sigma^2\rho}{\mathbb{E}_{\mu}[\rho^2\sigma^2]}, R_{\max}\right\}, \quad \text{where} \quad \alpha = \sqrt{2\log 2\,\frac{\mathbb{E}_{\mu}[\rho^2\sigma^2]}{n}}. \tag{24}$$

This setting satisfies the bound constraints by construction. A quick substitution also verifies that the second constraint is met. Consequently, it suffices to set ϵ to the value attained in the first constraint. Substituting the value of Δ in the constraint, we see that

$$\begin{aligned} \mathbb{E}_{\Delta}[\rho(x,a)r(x,a)] &= \mathbb{E}_{x\sim\lambda, a\sim\mu}[\rho(x,a)\Delta(x,a)] \\ &\geq \mathbb{E}_{x\sim\lambda,a\sim\mu}\left[\rho\frac{\alpha\sigma^2\rho}{\mathbb{E}_{\mu}[\rho^2\sigma^2]}\mathbf{1}(\rho\sigma^2\alpha \geq R_{\max}\mathbb{E}_{\mu}[\rho^2\sigma^2])\right] \\ &= \alpha\left(1 - \frac{\mathbb{E}_{\mu}[\rho^2\sigma^2\mathbf{1}(\rho\sigma^2\alpha \leq R_{\max}\mathbb{E}_{\mu}[\rho^2\sigma^2])]}{\mathbb{E}_{\mu}[\rho^2\sigma^2]}\right). \end{aligned}$$

Taking squares and dividing by 4 (since this is the lower bound on $2\sqrt{\epsilon}$ yields the statement of the theorem.

A.2 Proof of Theorem 4

We now give the proof of Theorem 4. This shares a lot of common ideas with the proof of Theorem 3, but with a crucial difference. In Theorem 3, there is non-trivial noise in the reward function, unlike in Theorem 4. This allowed the proof to work with just two candidate mean reward functions, since any realization in the data is corrupted with noise. However, in the absence of added noise, the task of mean identification becomes rather trivial: an estimator can just check whether η_1 or η_2 matches the observations exactly.

In order to avoid this pathology, we instead construct a richer family of reward functions which will be described next. In this construction, rather than just two mean rewards, the construction will involve a randomized design of the expected reward function from an appropriate prior distribution. The draw of the mean reward from a prior will essentially generate noise even though any given problem is noiseless. The construction will also highlight the crucial sources of difference between the contextual and multi-armed bandit problems, since the arguments rely on having a rich context distribution.

Creating a family of problems As before, we fix π, μ and λ and parametrize the problem distribution in terms of the mean reward function $\eta(x,a)$. However, now $\eta(x,a)$ is itself a random variable, which is drawn from a prior distribution. The distribution of $\eta(x,a)$ is a scaled Bernoulli, parametrized by a prior function $\theta(x,a)$ as described below:

$$\eta(x,a) = \begin{cases} R_{\max}(x,a) & \text{with probability } \theta(x,a) \\ 0 & \text{with probability } 1-\theta(x,a) \end{cases}. \tag{25}$$

We now set $D_\eta(r \mid x, a) = \eta(x, a)$. This clearly satisfies the constraints on the mean $0 \leq \mathbb{E}[r \mid x, a] \leq R_{\max}(x, a)$ by construction, and also $\text{Var}(r \mid x, a) = 0$ as per the setting of Theorem 4. The goal of an estimator is to take n samples generated according with $x \sim \lambda$, $a \mid x \sim \mu$ and $r \mid x, a \sim D_\eta$ and output an estimate $\hat{v}$ such that $\mathbb{E}_\eta[(\hat{v} - v_\eta^\pi)^2]$ is small. We recall our earlier shorthand v_η to denote the value of π under the reward distribution generated by η. For showing a lower bound on this quantity, it is clearly sufficient to pick any prior distribution governed by a parameter θ as in Equation 25, and instead lower bound $\mathbb{E}_\theta\left[\mathbb{E}_\eta[(\hat{v} - v_\eta)^2 \mid \eta]\right]$. If the expectation is large for an estimator $\hat{v}$, then there must be at least one function $\eta(x, a)$ which induces a large error for $\hat{v}$ which is the goal of the lower bound. Consequently, we focus in the proof on lower bounding this quantity, which can be decomposed by Cauchy-Shwartz inequality into:

$$\mathbb{E}_\theta\left[\mathbb{E}_\eta[(\hat{v} - v_\eta)^2 \mid \eta]\right] \geq \frac{1}{2}\mathbb{E}_\theta\left[\mathbb{E}_\eta[(\hat{v} - \mathbb{E}_\theta[v_\eta])^2 \mid \eta]\right] - \mathbb{E}_\theta\left[(v_\eta - \mathbb{E}_\theta[v_\eta])^2\right].$$

Then we can further take a worst-case over all problems in the above inequality to obtain

$$\begin{aligned}\sup_\eta \mathbb{E}_\eta[(\hat{v} - v_\eta)^2] &\geq \sup_\theta \mathbb{E}_\theta\left[\mathbb{E}_\eta[(\hat{v} - v_\eta)^2 \mid \eta]\right] \\ &\geq \underbrace{\sup_\theta \frac{1}{2}\mathbb{E}_\theta\left[\mathbb{E}_\eta[(\hat{v} - \mathbb{E}_\theta[v_\eta])^2 \mid \eta]\right]}_{\mathcal{T}_1} - \underbrace{\mathbb{E}_\theta\left[(v_\eta - \mathbb{E}_\theta[v_\eta])^2\right]}_{\mathcal{T}_2}. \end{aligned} \tag{26}$$

This decomposition is rather interesting. It says that the expected MSE of an estimator in estimating v_η can be related to the MSE of the same estimator in estimating the quantity $\mathbb{E}_\theta[v_\eta]$, as long as the variance of the quantity v_η under the distribution generated by θ is not too large. This is a very important observation, since we can now choose to instead study the MSE of an estimator in estimating $\mathbb{E}_\theta[v_\eta]$ as captured by $\mathcal{T}_1$. Unlike the distribution D_η which is degenerate, this problem has a non-trivial noise arising from the randomized draw of η according to θ. Thus we can use similar techniques as the proof of Theorem 4, albeit where the reward distribution is a scaled Bernoulli instead of Gaussian. For now, we focus on controlling $\mathcal{T}_1$, and $\mathcal{T}_2$ will be handled later.

In order to bound $\mathcal{T}_1$, we will consider two carefully designed choices θ_1 and θ_2 to induce two different problem instances and show that $\mathcal{T}_1$ is large for *any estimator* under one of the two parameters. In doing this, it would be convenient to use the additional shorthand $\ell_\theta(\hat{v}) = (\hat{v} - \mathbb{E}_\theta[v_\eta])^2$. Proceeding as in the proof of Theorem 3, we have

$$\begin{aligned}\mathcal{T}_1 &= \sup_\theta \mathbb{E}_\theta\left[\mathbb{E}_\eta[(\hat{v} - \mathbb{E}_\theta[v_\eta])^2 \mid \eta]\right] \geq \sup_\theta \mathbb{E}_\theta\left[\mathbb{E}_\eta[\ell_\theta(\hat{v}) \mid \eta]\right] \\ &\geq \epsilon \sup_\theta \mathbb{P}_\theta\left(\ell_\theta(\hat{v}) \geq \epsilon\right) \geq \epsilon \max_{\theta \in \{\theta_1, \theta_2\}} \mathbb{P}_\theta\left(\ell_\theta(\hat{v}) \geq \epsilon\right) \\ &\geq \frac{\epsilon}{2}\left[\mathbb{P}_{\theta_1}\left(\ell_{\theta_1}(\hat{v}) \geq \epsilon\right) + \mathbb{P}_{\theta_2}\left(\ell_{\theta_2}(\hat{v}) \geq \epsilon\right)\right].\end{aligned}$$

Reduction to hypothesis testing As in the proof of Theorem 3, we now reduce the estimation problem into a hypothesis test for whether the data was generating according to the parameters θ_1 or θ_2. The arguments here are similar to the earlier proof, and we will be a bit terse in this presentation.

As before, our hypothesis test has entire knowledge of D_η as well as θ_1 and θ_2. Consequently, we construct a test based on picking θ_1 whenever $\ell_{\theta_1}(\hat{v}) \leq \ell_{\theta_2}(\hat{v})$. As before, we will ensure that $|\mathbb{E}_{\theta_1}[v_\eta] - \mathbb{E}_{\theta_2}[v_\eta]| \geq 2\sqrt{\epsilon}$ so that for any estimator $\hat{v}$, we have

$$\ell_{\theta_1}(\hat{v}) + \ell_{\theta_2}(\hat{v}) \geq 2\epsilon.$$

By a similar argument, we further conclude that the error of our hypothesis test under a uniform prior on the choices of θ is at most

$$\frac{1}{2}\left[\mathbb{P}_{\theta_1}\left(\ell_{\theta_1}(\hat{v}) \geq \epsilon\right) + \mathbb{P}_{\theta_2}\left(\ell_{\theta_2}(\hat{v}) \geq \epsilon\right)\right]$$

Invoking Le Cam's argument Once again, we can lower bound the error rate of our test by invoking the result of Le Cam. Doing so, requires an upper bound on the KL-divergence $D_{KL}(P_{\theta_1} \| P_{\theta_2})$. The only difference from our earlier argument is that these distributions are now Bernoullis instead of Normals based on the construction in Equation 25. More formally, we have

$$\begin{aligned} D_{KL}(P_{\theta_1} \| P_{\theta_2}) &= \int \log \frac{p(r; \theta_1(x,a))}{p(r; \theta_2(x,a))} dp(r; \theta_1(x,a))\, dp(x)\, d\mu(a|x) \\ &= \mathbb{E} D_{KL}(\text{Ber}(\theta_1(x,a)) \| \text{Ber}(\theta_2(x,a)))). \end{aligned}$$

Invoking Lemma 4, we can further upper bound this as

$$D_{KL}(P_{\theta_1} \| P_{\theta_2}) \leq \frac{1}{4}\mathbb{E}_\mu\left[(\theta_1(x,a) - \theta_2(x,a))^2\right].$$

Picking the parameters It remains to carefully choose θ_1 and θ_2. We define θ_2 to be a uniform prior: $\theta_2(x,a) \equiv 0.5$, and let $\theta_1(x,a) = \theta_2(x,a) + \Delta(x,a)$, where $\Delta(x,a)$ will be chosen to satisfy certain constraints as before. Collecting the constraints, we get the maximization problem

$$\begin{aligned} \max_{\Delta} \quad & \epsilon \\ \text{such that} \quad & \mathbb{E}_\mu[\rho(x,a)\Delta(x,a)R_{\max}(x,a)] \geq 2\sqrt{\epsilon}, \\ & \frac{1}{4}\mathbb{E}\left[\Delta(x,a)^2\right] \leq \frac{\log 2}{n} \quad \text{and} \\ & 0 \leq \Delta(x,a) \leq 0.5. \end{aligned}$$

For some $\alpha > 0$ to be determined shortly, we pick

$$\Delta = \min\left\{\frac{\rho R_{\max}\alpha}{\mathbb{E}_\mu \rho^2 R_{\max}^2}, 0.5\right\}.$$

The bound constraints are satisfied by construction and we set $\alpha = \sqrt{4\log 2\, \mathbb{E}_\mu \rho^2 R_{\max}^2 / n}$ to satisfy the second constraint. A feasible choice of ϵ then yields

$$\begin{aligned} 2\sqrt{\epsilon} = \mathbb{E}_\mu[\rho(x,a)\Delta(x,a)R_{\max}(x,a)] &\geq \mathbb{E}_\mu\left[\frac{\rho^2 R_{\max}^2 \alpha \mathbf{1}(\rho R_{\max} \leq \mathbb{E}_\mu[\rho^2 R_{\max}^2]/2\alpha)}{\mathbb{E}_\mu[\rho^2\alpha^2]}\right] \\ &= \alpha - \alpha\frac{\mathbb{E}_\mu\left[\rho^2 R_{\max}^2 \mathbf{1}(\rho R_{\max} \leq \mathbb{E}_\mu[\rho^2 R_{\max}^2]/2\alpha)\right]}{\mathbb{E}_\mu[\rho^2 R_{\max}^2]}. \end{aligned}$$

Collecting our arguments this far, we have established that

$$
\begin{aligned}
\mathcal{T}_1 &\geq \frac{\epsilon}{2}\left[\mathbb{P}_{\theta_1}\left(\ell_{\theta_1}(\hat{v}) \geq \epsilon\right) + \mathbb{P}_{\theta_2}\left(\ell_{\theta_2}(\hat{v}) \geq \epsilon\right)\right] \\
&\geq \frac{\epsilon}{2}\exp(-nD_{KL}(P_{\theta_1}\|P_{\theta_2})) \geq \frac{\epsilon}{4} \\
&\geq \frac{\log 2\mathbb{E}_\mu \rho^2 R_{\max}^2}{4n}\left[1 - \frac{\mathbb{E}_\mu\left(\rho^2 R_{\max}^2 \mathbf{1}\left(\left\{\rho R_{\max} > \sqrt{n\mathbb{E}_\mu(\rho^2\mathbb{R}_{\max}^2)/(16\log 2)}\right\}\right)\right)}{\mathbb{E}_\mu \rho^2 R_{\max}^2}\right]^2
\end{aligned}
$$

In order to complete the proof, we need to further upper bound $\mathcal{T}_2$ in the decomposition (26). This is done in Lemma 2, and plugging this upper bound in Equation 26 completes the proof.

Lemma 2. *Let $\lambda(x)$ be any context distribution. Then for any distribution θ over the random variable $\eta(x,a) = \mathbb{E}[r \mid x, a]$, satisfying $0 \leq \eta(x,a) \leq R_{\max}(x,a)$, we have*

$$
\mathbb{E}_\theta\left[(v_\eta - \mathbb{E}_\theta[v_\eta])^2\right] \leq \lambda_0 \log(4/\lambda_0)\mathbb{E}_\mu[R_{\max}^2\rho^2] \ ,
$$

where $\lambda_0 = \max_{x\in\mathcal{X}} \lambda(x)$.

Proof. We start by bounding the range of $(v_\eta - \mathbb{E}_\theta[v_\eta])^2$, viewed as a random variable under θ. From the definition of v_η,

$$
0 \leq v_\eta \leq \mathbb{E}_\pi[R_{\max}] = \mathbb{E}_\mu[\rho R_{\max}] \ ,
$$

so also $0 \leq \mathbb{E}_\theta[v_\eta] \leq \mathbb{E}_\mu[\rho R_{\max}]$. Hence, $|v_\eta - \mathbb{E}_\theta[v_\eta]| \leq \mathbb{E}_\mu[\rho R_{\max}]$, and we obtain the bound

$$
(v_\eta - \mathbb{E}_\theta[v_\eta])^2 \leq (\mathbb{E}_\mu[\rho R_{\max}])^2 \leq \mathbb{E}_\mu[\rho^2 R_{\max}^2], \tag{27}
$$

because $\rho R_{\max} \geq 0$.

The proof proceeds by applying Hoeffding's inequality to control the probability that $(v_\eta - \mathbb{E}_\theta[v_\eta])^2 \geq t^2$ for a suitable t. Then we can, with high probability, use the bound $(v_\eta - \mathbb{E}_\theta[v_\eta])^2 \geq t^2$, and with the remaining small probability apply the bound of Eq. (27).

To apply Hoeffding's inequality, we write v_η explicitly as

$$
v_\eta = \sum_{x\in\mathcal{X}}\sum_{a\in\mathcal{A}} \eta(x,a)\rho(x,a)\mu(a|x)\lambda(x) =: \sum_{x\in\mathcal{X}}\sum_{a\in\mathcal{A}} Y(x,a).
$$

That is, the quantity v_η can be written as a sum of a finite number of independent random variables, where we use the fact that the number of contexts is finite. Note that $\mu(a|x) \leq 1$ and $\lambda(x) \leq \lambda_0$, so each summand obeys

$$
0 \leq Y(x,a) \leq R_{\max}(x,a)\rho(x,a)\sqrt{\mu(a|x)\lambda(x)}\sqrt{\lambda(x)} \leq R_{\max}(x,a)\rho(x,a)\sqrt{\mu(a|x)\lambda(x)\lambda_0} \ .
$$

By Hoeffding's inequality, we have

$$
\begin{aligned}
\mathbb{P}(|v_\eta - \mathbb{E}_\theta v_\eta| \geq t) &\leq 2\exp\left\{-\frac{2t^2}{\sum_{x\in\mathcal{X}}\sum_{a\in\mathcal{A}} R_{\max}^2(x,a)\rho^2(x,a)\mu(a|x)\lambda(x)\lambda_0}\right\} \\
&= 2\exp\left\{-\frac{2t^2}{\lambda_0\mathbb{E}_\mu[R_{\max}^2\rho^2]}\right\}.
\end{aligned}
$$

Take $t = \sqrt{\lambda_0 \log(4/\lambda_0)\mathbb{E}_\mu[R_{\max}^2\rho^2]/2}$ in the above bound, which yields

$$
\mathbb{P}\Big[(v_\eta - \mathbb{E}_\theta v_\eta)^2 \geq t^2\Big] = \mathbb{P}\Big[|v_\eta - \mathbb{E}_\theta v_\eta| \geq t\Big] \leq \frac{\lambda_0}{2} \ .
$$

Thus, using Eq. (27), we have

$$\begin{aligned}\mathbb{E}_\theta\left[(v_\eta - \mathbb{E}_\theta v_\eta)^2\right] &\le t^2 \mathbb{P}\Big[(v_\eta - \mathbb{E}_\theta v_\eta)^2 < t^2\Big] + \mathbb{E}_\mu[\rho^2 R_{\max}^2] \cdot \mathbb{P}\Big[(v_\eta - \mathbb{E}_\theta v_\eta)^2 \ge t^2\Big] \\ &\le t^2 + \mathbb{E}_\mu[\rho^2 R_{\max}^2] \cdot \frac{\lambda_0}{2} \\ &= \frac{\lambda_0 \log(4/\lambda_0) \mathbb{E}_\mu[R_{\max}^2 \rho^2]}{2} + \mathbb{E}_\mu[\rho^2 R_{\max}^2] \cdot \frac{\lambda_0}{2} \\ &\le \lambda_0 \log(4/\lambda_0) \mathbb{E}_\mu[R_{\max}^2 \rho^2] \ .\end{aligned}$$

□

B Other proofs

Proof of Lemma 1. The variance (therefore MSE) of the importance weighting estimator can be decomposed into two parts

$$\begin{aligned}\mathrm{Var}\left(\hat{v}_{\mathrm{IPS}}\right) &= \mathbb{E}_\mu(\mathrm{Var}(\hat{v}_{\mathrm{IPS}}|x,a) + \mathrm{Var}_\mu(\mathbb{E}(\hat{v}_{\mathrm{IPS}}|x,a)) = \frac{1}{n}\mathbb{E}_\mu[\rho^2 \mathrm{Var}(r|x,a)] + \frac{1}{n}\mathrm{Var}_\mu[\rho \mathbb{E}(r|x,a)] \\ &\le \frac{1}{n}\mathbb{E}_\mu[\rho^2\sigma^2] + \frac{1}{n}\mathbb{E}_\mu[\rho^2 \mathbb{E}(r|x,a)^2] \le \frac{1}{n}\mathbb{E}_\mu[\rho^2\sigma^2] + \frac{1}{n}\mathbb{E}_\mu[\rho^2 R_{\max}^2]\end{aligned}$$

□

Proof of Theorem 2. Denote $A_x \subseteq \mathcal{A}$ to be $A_x = \left\{a \in \mathcal{A} \middle| \frac{\pi(a|x)}{\mu(a|x)} \le \tau\right\}$. For brevity, we also denote $A_i := A_{x_i}$. We decompose the mean square error and control each term separately.

$$\mathrm{MSE}(\hat{v}_{\mathrm{SWITCH}}) = |\mathbb{E}\hat{v}_{\mathrm{SWITCH}} - v^\pi|^2 + \mathrm{Var}(\hat{v}_{\mathrm{SWITCH}}).$$

We first calculate the bias. Note that bias is incurred only in the terms that fall in A_x^C, so that

$$\begin{aligned}\mathbb{E}\hat{v}_{\mathrm{SWITCH}} - v^\pi &= \mathbb{E}\left[\sum_{a \in A_x^c} \hat{r}(x,a)\pi(a|x)\right] - \mathbb{E}\left[\sum_{a \in A_x^c} \mathbb{E}(r|x,a)\pi(a|x)\right] \\ &= \mathbb{E}_\pi\left[(\hat{r}(x,a) - \mathbb{E}(r|x,a))\mathbf{1}(a \in A_x^c)\right] = \mathbb{E}_\pi\left[\hat{r}(x,a) - \mathbb{E}(r|x,a)|\ a \in A_x^c\right]\mathbb{P}_\pi(a \in A_x^c) \\ &= \mathbb{E}_\pi[\epsilon(x,a)|\ a \in A_x^c]\mathbb{P}_\pi(a \in A_x^c),\end{aligned}$$

where we recall that $\epsilon(x,a) = \hat{r}(x,a) - \mathbb{E}[r|\ x,a]$.

Next we upper bound the variance. Note that the variance contributions from the IPS part and the DM part are not independent, since the indicators $\rho(x_i,a) > \tau$ and $\rho(x_i,a) \le \tau$ are mutually exclusive. To simplify the analysis, we use the following inequality that holds for any random variable X and Y:

$$\mathrm{Var}(X+Y) \le 2\mathrm{Var}(X) + 2\mathrm{Var}(Y).$$

This allows us to calculate the variance of each part separately.

$$\begin{aligned}
\mathrm{Var}(\hat{v}_{\text{SWITCH}}) \leq & 2\,\mathrm{Var}\left(\frac{1}{n}\sum_{i=1}^{n}[r_i\rho_i\mathbf{1}(a_i \in A_i)]\right) + 2\,\mathrm{Var}\left(\frac{1}{n}\sum_{i=1}^{n}\sum_{a\in\mathcal{A}}\hat{r}(x,a)\pi(a|x)\mathbf{1}(a \in A_x^C)\right) \\
= & \frac{2}{n}\,\mathrm{Var}_\mu\left[r\rho\mathbf{1}(a \in A_x)\right] + \frac{2}{n}\,\mathrm{Var}\left(\sum_{a\in A_x^c}\hat{r}(x,a)\pi(a|x)\right) \\
= & \frac{2}{n}\,\mathbb{E}_\mu\mathrm{Var}\left[r\rho\mathbf{1}(a \in A_x)|\, x,a\right] + \frac{2}{n}\,\mathrm{Var}_\mu\mathbb{E}[r\rho\mathbf{1}(a \in A_x)|\, x,a] + \frac{2}{n}\,\mathrm{Var}\left(\sum_{a\in A_x^c}\hat{r}(x,a)\pi(a|x)\right) \\
\leq & \frac{2}{n}\,\mathbb{E}_\mu\mathrm{Var}\left[r\rho\mathbf{1}(a \in A_x)|\, x,a\right] + \frac{2}{n}\,\mathbb{E}_\mu(\mathbb{E}[r\rho\mathbf{1}(a \in A_x)|\, x,a])^2 + \frac{2}{n}\,\mathbb{E}\left(\sum_{a\in A_x^c}\hat{r}(x,a)\pi(a|x)\right)^2 \\
\leq & \frac{2}{n}\,\mathbb{E}_\mu\left[\sigma^2\rho^2\mathbf{1}(a \in A_x)\right] + \frac{2}{n}\,\mathbb{E}_\mu\left[R_{\max}^2\rho^2\mathbf{1}(a \in A_x)\right] + \frac{2}{n}\,\mathbb{E}\left(\sum_{a\in A^c}\hat{r}(x,a)\pi(a|x)\right)^2.
\end{aligned}$$

To complete the proof, note that the first term is further upper bounded using Jensen's inequality as

$$\begin{aligned}
\mathbb{E}\left(\sum_{a\in A_x^c}\hat{r}(x,a)\pi(a|x)\right)^2 &= \mathbb{E}\left[\left(\sum_{a\in A_x^C}\pi(a|x)\right)^2\left(\sum_{a\in A_i^c}\frac{\hat{r}(x,a)\pi(a|x)}{\sum_{a\in A_x^C}\pi(a|x)}\right)^2\right] \\
&\leq \mathbb{E}\left[\left(\sum_{a\in A_x^C}\pi(a|x)\right)\left(\sum_{a\in A_x^C}\hat{r}(x,a)^2\pi(a|x)\right)\right] \\
&\leq \mathbb{E}_\pi\left[R_{\max}^2\mathbf{1}(\rho > \tau)\right],
\end{aligned}$$

where the final inequality uses $\sum_{a\in A_x^C}\pi(a|x) \leq 1$ and $\hat{r}(x,a) \in [0, R_{\max}(x,a)]$ almost surely.

Combining the bias and variance bounds, we get the stated MSE upper bound. □

C Utility Lemmas

Lemma 3 (Hoeffding, 1963, Theorem 2)**.** *Let $X_i \in [a_i, b_i]$ and $X_1, ..., X_n$ are drawn independently. Then the empirical mean $\bar{X} = \frac{1}{n}(X_1 + ... + X_n)$ obeys*

$$\mathbb{P}(|\bar{X} - \mathbb{E}[\bar{X}]| \geq t) \leq 2e^{-\frac{2n^2t^2}{\sum_{i=1}^n(b_i-a_i)^2}}.$$

Lemma 4 (Bernoulli KL-divergence)**.** *For $0 < p, q < 1$, we have*

$$D_{\mathrm{KL}}(\mathrm{Ber}(p)\|\mathrm{Ber}(q)) \leq (p-q)^2(\frac{1}{q} + \frac{1}{1-q}).$$

Proof.

$$\begin{aligned}
D_{\mathrm{KL}}(\mathrm{Ber}(p)\|\mathrm{Ber}(q)) &= p\log(\frac{p}{q}) + (1-p)\log(\frac{1-p}{1-q}) \\
&\leq p\frac{p-q}{q} + (1-p)\frac{q-p}{1-q} = \frac{(p-q)^2}{q} + (p-q) + \frac{(p-q)^2}{1-q} + (q-p) \\
&= (p-q)^2(\frac{1}{q} + \frac{1}{1-q}).
\end{aligned}$$

□

References

Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

Oliver Bembom and Mark J van der Laan. Data-adaptive selection of the truncation level for inverse-probability-of-treatment-weighted estimators. 2008.

Léon Bottou, Jonas Peters, Joaquin Quinonero Candela, Denis Xavier Charles, Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Y Simard, and Ed Snelson. Counterfactual reasoning and learning systems: the example of computational advertising. *Journal of Machine Learning Research*, 14(1): 3207–3260, 2013.

Claes M Cassel, Carl E Särndal, and Jan H Wretman. Some results on generalized difference estimation and generalized regression estimation for finite populations. *Biometrika*, 63(3):615–620, 1976.

Corinna Cortes, Yishay Mansour, and Mehryar Mohri. Learning bounds for importance weighting. In *Advances in neural information processing systems*, pages 442–450, 2010.

Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *ICML*, 2011.

Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4):485–511, 2014.

Arthur Gretton, Alex Smola, Jiayuan Huang, Marcel Schmittfull, Karsten Borgwardt, and Bernhard Schölkopf. Covariate shift by kernel mean matching. *Dataset shift in machine learning*, 3(4):5, 2009.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.

Paul W Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396): 945–960, 1986.

Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.

Nan Jiang and Lihong Li. Doubly robust off-policy evaluation for reinforcement learning. In *ICML'16*, 2016.

John Lafferty, Han Liu, and Larry Wasserman. Minimax theory, 2008. URL http://www.stat.cmu.edu/~larry/=sml/Minimax.pdf.

Lihong Li, Wei Chu, John Langford, and Xuanhui Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 297–306. ACM, 2011.

Lihong Li, Rémi Munos, and Csaba Szepesvári. Toward minimax off-policy value estimation. In *AISTATS*, 2015.

Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*. John Wiley & Sons, 2002.

H Lock Oh and Frederick J Scheuren. Weighting adjustment for unit nonresponse. *Incomplete data in sample surveys*, 2:143–184, 1983.

James M Robins and Andrea Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.

Andrea Rotnitzky, Quanhong Lei, Mariela Sued, and James M Robins. Improved double-robust estimation in missing data and causal inference models. *Biometrika*, 99(2):439–456, 2012.

Yeneng Sun. The exact law of large numbers via fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126(1):31–69, 2006.

Philip S Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *ICML'16*, 2016.